

Childcare and the New Part-Time: Gender Gaps in Long-Hour Professions*

Kendall Houghton[†]

Job Market Paper

(latest version)

November 2, 2020

Abstract

The gender wage gap in the United States is persistent and especially pronounced at the top of the distribution. Recent worker surveys suggests this gap is in part driven by a difference in average work hours, even between men and women employed full-time. This paper examines gender differences in work timing and elasticity using hourly data on worker activity for approximately 145,000 tech workers in order to understand what leads to this difference in the number of hours worked. I find both genders work outside the traditional work week, but men work more than women on nights and weekends — times when formal childcare is relatively scarce. In order to isolate the impact of childcare, I examine how work activity varies in response to unexpected winter weather public school closures. Women respond to these unexpected breaks in childcare by reducing work activity by 34%. Male work activity does not respond to these unexpected breaks. The results provide evidence for the emerging theory that the persistent wage gap between men and women in high-wage professions is a function of a long-hour wage premium and the allocation of within-household childcare to women.

Keywords: Economics of gender, time allocation and labor supply, child care (JEL J13, J16, J22)

*I am deeply grateful to my co-chairs, Benjamin Hansen and Caroline Weber, and my committee members, Grant McDermott and Laura Leete. All errors are my own.

[†]Department of Economics, University of Oregon. Email: khoughto@uoregon.edu.

Introduction

The median female worker in the United States makes approximately \$10,000 less than the median male worker in the United States (U.S. Census Bureau, 2017a). Over the course of a lifetime, this difference in earnings leads to a significant wealth gap. Women hold 65% of U.S. student loan debt¹ (American Association of University Women, 2019) and are more likely to live below the poverty line at every age² (Joint Economic Committee, 2016). A large body of work seeks to explain the persistent earnings gap (for a thorough review, see Blau and Katz, 2017), and finds that a difference in weekly work hours can explain much of the difference.

Historically, the difference in work hours was largely a difference in part-time versus full-time employment. As technology has made it possible to work remotely, workers have been faced with the both the possibility and, in some cases expectation, that they work much more than traditional full-time. This phenomenon has been documented in the press (Miller, 2019; Miller, 2015a; Miller, 2015b), and described in surveys of the workforce (Cha and Weeden, 2014; Noonan, Corcoran, and Courant, 2005). This high-hour equilibrium is present in the top of the distribution and in professions that reward wage as a strongly nonlinear function of work hours (Goldin & Katz, 2016).

Female MBAs and JDs self-report working more than 40 hours a week, but less than the hours self-reported by men in the same samples (Noonan et al., 2005, Bertrand, Goldin, and Katz, 2010). Self-reports from workers across all fields indicate this gender difference is a broad phenomenon (Cha & Weeden, 2014). In this paper, I use minute-level records of computer programmer work activity to verify the presence of long-work hours in STEM professions, and that women are working fewer hours than men. I then consider the mechanism behind this behavior. I document precisely when during the week and day men are

¹For the past 18 years, women have earned 57% of bachelor degrees in the U.S. (National Center for Education Statistics, 2017).

²For individuals aged 65 to 74, 9.8% of women are in poverty and 7.7% of men. For individuals above the age of 75, 8% of men and 12.1% of women live in poverty, although this statistic is confounded by more women living later in life than men (U.S. Census Bureau, 2017b).

participating in formal labor and women are not, and find that women are less likely than men to work on nights and weekends.

Women may have a greater taste for leisure than men, and cultural factors may make leisure more rewarding during these hours. On the other hand, these are also times when formal childcare is less readily available. Female labor force participation is inextricably connected to childcare availability (Blau & Kahn, 2013). We know that the percent of mothers who work more than forty hours a week is significantly below the percent in other demographic groups, and the gap between the share of mothers and the share of non-mothers participating in these hours has been increasing since the 1990's (Weeden, Cha, and Bucca, 2016; Cha and Weeden, 2014; Noonan et al., 2005). This is highly suggestive that childcare is the factor that limits mothers from working the hours that women without children and men are working.

In order to separate childcare availability from other potential mechanisms, I consider shocks to childcare availability. In my estimation, I exploit winter weather public school closures as a natural experiment in which childcare is unexpectedly unavailable. I consider female and male work activity on these days.

Male work activity does not respond to unexpected shocks in the public school calendar, but female work activity does. Female activity decrease by 34% on these days. These results are consistent with a story that households plan for childcare external to the household, but that women are more likely to be responsible for within-household childcare when external childcare is unavailable.

It is possible that the parental status of the women in my data is systematically different than the parental status of the men in my data. I could, for example, be comparing single men to women with children. In order to increase the likelihood that I am comparing men with children to women with children, I augment my primary data with estimated age for a subset of individuals. I use publicly available information from LinkedIn to estimate the

individual’s age using their education levels and corresponding graduation years³. I match individuals by limiting my LinkedIn search to those working in STEM fields, in the geographic location listed in my primary dataset, with a matching name and company affiliation. When age is included, the average effect for women disappears, and a much larger effect for women between the age of 30 and 40 emerges.

The wage gap is widest among college graduates and at top of the wage distribution (Goldin, Kerr, Olivetti, & Barth, 2017). Science, technology, engineering, and math (STEM) are often highlighted as fields in which increased gender parity is wanted, but increases in the number of female college students enrolled in these majors have not been reflected in the workforce (Speer, 2019). Female STEM graduates are less likely to take STEM jobs, and are more likely to leave the STEM workforce. Women who stay in STEM advance more slowly, and are less likely to hold management positions. These gaps have been resistant to intervention (van den Hurk, Meelissen, & van Langen, 2019). A lessened ability to work long hours may be the cause of the slow advancement and lower participation rate. My work suggests that in STEM careers, and perhaps other high-hour professions, the wage gap is not only an issue of childcare during traditional work hours, but all the other times as well.

1 Data

GitHub. In order to document worker activity at a precise timescale, I use a publicly available record of worker activity from GitHub, a version-control platform with more than 40 million users (GitHub, 2019). Workers who use technology use version control to manage solo and collaborative projects. A basic example of version control involves saving multiple versions of a file to a local hard drive. Consider a common example of this practice, for *File*. A user might save this file as *File1*, *File2*, *File3*, etc. as both minor and major changes are made to the file. For complex projects, version control software provides a more user-friendly

³“Scraping” data from LinkedIn in this manner has been ruled legal by federal courts (United States Court of Appeals for the Ninth Circuit, 2019)

version of this practice. The worker first creates one file. After making changes to this file, the worker saves these changes using version control. These changes, and the lines of code or text within them, are tracked by the version control software in order to allow the worker to return to any previous version at any point. Crucially for this project, these changes as well as other activity on the version control platform are tracked by time and date. I am able to exploit the time-date system of tracking to create a time-series of worker activity.

In GitHub, activity is referred to as an event. For a detailed list and description of these events see Table A2. Table A1 summarizes how frequently each type of activity occurs. The second most common type of event is a “push” event, which is when a user takes file changes from their local computer to the remote copy of the file. Inside the push event details, I am able to observe each time the user has made a change to that file on their local machine. These local changes are “commits”, the most common event type, and I observe portions of their content and the timestamp of when the file was saved. Together, commits and push events represent 70% of the data. In my analysis, I limit my sample to these events. Figures A1 and A2, show activity across the day for each event type. Figure A1 displays this for each event type, and Figure A2 displays this for the ten most common events. All events outside the top ten represent less than 1% of the data. The data pattern for all types is qualitatively similar, with less noise as event type frequency increases.

I access the record of GitHub activity from two sources, GitHub Archive and GitHub Torrent. GitHub Archive is log of all public events that happen on GitHub. This data source is updated hourly and covers activity from 2011 forward. GitHub Torrent is meant to capture the relationships between the different aspects of GitHub – the relationships between users, files and users, and files with each other. This dataset records the stated location of users as well. GitHub torrent is updated monthly and covers 2013 to present.

I gather user name and event activity for all events except for commits from GitHub Archive. I use GitHub Torrent to collect the user location and commit activity. I merge these datasets together using the login name of the user. The raw data is processed minimally.

The timestamp of activity is standardized to Universal Coordinated Time in both GitHub Archive and GitHub Torrent. In order to track when the user is working during the day, I adjust timestamps by the timezone provided in their geographic information. Figure A3 shows the daily activity of users by timezone after this transformation. Each timezone follows a similar daily pattern.

I use two subsamples of the worker data for my analysis. When I examine general trends for men and women, I consider push and commit activity between 2017 and 2019 for users with a stated location in the United States and an identifiable first name. This sample is 145,333 users. When I consider the effect of school day closures, I limit my analysis to Seattle, Washington. There are 4,392 users in this smaller sample.

Gender Imputation. I impute gender for each user. In the GitHub data, I am able to identify the first name of users who provide a first name in their profile information. In order to impute gender for each user, I use the the R package, ‘gender’, which uses Census data to predict the gender of an individual when given their first name and geographic location (Blevins & Mullen, 2015). I provide the first name of the user as collected from GitHub Archive and the geographic location as collected from GitHub Torrent.

School Snow Closure Records. In my analysis of winter weather school closures, I use a subsample of users in Seattle, Washington. I use the Twitter page of Seattle Public School District to identify snow and winter weather school closure dates. Most school districts do not maintain public, formal records of closures due to snow. Fortunately, Seattle Public School District maintains an active Twitter page that communicates news of school closures, delays, and changes to activity schedules. The tweets announcing these decisions remain on the Twitter page for the district. I locate relevant tweets on the Seattle Public School District page and verify these dates by searching for tweets by other users on the same day documenting the school closure and winter weather conditions.

LinkedIn. As part of my robustness analysis, I impute the age of users to identify likely parents. Through the GitHub data, I am able to identify user first name, last name, and company affiliations. With this information, I identify these users in LinkedIn using a web-based LinkedIn URL finder application, *PhantomBuster*. From the LinkedIn profiles, I scrape user education and dates of graduation using Selenium⁴. Using this information, I impute a likely age for the users.⁵

2 Overall Work Patterns

I first document the typical behavior of all users in my data. Given the high frequency of the data, it is possible to track the daily work schedule of the average user. In Figure 1, I plot the total amount of activity that occurs for all users on any given day of the week for every hour of that day. I scale this total amount of activity by dividing by the max total amount of activity observed for a day-of-week, hour-of-day combination.

On average, users are most active between 2pm and 3pm on Wednesday and Thursday. Activity is lowest between the hours of 1am and 7am across all days. On workdays, activity increases most sharply between 8am and 9am. There is a dip in activity between 11am and 1 o'clock that is consistent with the lunch hour. Activity decreases most sharply between 5pm and 6pm. This decrease is most pronounced on Friday.

Evenings follow a similar pattern across the week. Monday, Tuesday, Wednesday, Thursday, and Sunday nights are closely related as are Friday and Saturday nights. These groupings correspond to nights before work days and nights before weekend days. Monday and Tuesday night have the most activity, followed by Wednesday, then Thursday, and then Sunday. For the nights that occur on a work day and before a work day, there is a decrease in activity as the week progresses. After 5pm, the most work occurs at 9pm, regardless of day. Friday

⁴For more information on the scraping process and code, see my co-authored work on COVID-19 with Ben Hansen, Grant McDermott, and Caroline Weber

⁵I assume individuals graduate with a Bachelor's degree at age 22, with a Master's degree at age 25, and with a PhD at age 30. The earliest listed degree and graduation date is used, as deviations from the average education path timing become more likely over time.

and Saturday are very similar from 8pm onward.

3 Description of Male and Female Work Habits

I begin the analysis of gender difference by examining the differences in male and female work timing across the day and week.

3.1 Empirical Method

For my analysis, I use the following general specification, which compares the activity for men and women over some time scale. For the following equation, I am using the example of comparing male and female activity across the work week.

$$\text{Log}(\text{Activity}_{tgd}) = \beta_0 + \theta \text{Female}_g + \sum_{d=1}^7 \delta_d \text{Day}_t + \sum_{d=1}^7 \gamma_d \text{Female}_g \times \text{Day}_t + \epsilon_{tgd}$$

Activity is a general term that refers to the combination of push and commit events as described in Section 1. This is the total amount of activity on date t for users of gender g on day of the week d . I include a fixed effect for activity by users identified as female, *Female*, in order to control for baseline differences in activity levels. In each specification, I consider various fixed effects, including month of year. In other specifications, these day of week effects are included as controls. I allow for errors with heteroskedasticity and autocorrelation by using Newey-West estimator and allowing for a two week lag. I am interested in the coefficients attached to the interaction between day of week and the female indicator variable, γ_d , which capture the differential day of week effects for women.

3.2 Results

Men and women display different daily and weekly work timing, on average. Figure 2 illustrates the difference in activity by hour of day during the work week, and Figure 3

illustrates the difference in activity by hour of day during the weekend. In Figure 4, I plot activity by day of week. Men tend to have a more diffuse work schedule, while female users are more active during times associated with a traditional work schedule.

Work Week. During the week, men and women both have the most activity during traditional 9am to 5pm work hours. However, the difference between work during this time and other hours is much more pronounced for women. The precise estimates for these differences are reported in Table 1. Between 6pm and 2am, female work activity declines more than male work activity for every hour, but the magnitude of the difference varies. The largest difference occurs between 6 and 7pm. The smallest difference occurs between 8 and 10pm. These details are shown in Figure 2.

During the traditional work day, women have less variable work activity. There is no statistically significant difference between men and women during 9am to 3pm. Women allocate their work to the 4pm to 5pm period of time 11% more than men do, but as illustrated in Figure 2, this is because male work decreases during these hours and female work does not.

In Section 4, I explore how child care responsibilities contribute to differences in male and female work activity. The large difference between 6pm and 7pm with the lessened difference between 8pm and 10pm is consistent with a story of women having non-work responsibilities in the household while children are awake and during traditional dinner preparation hours.

Weekends. On weekends, male work activity is allocated in a pattern that closely resembles the weekday work pattern. As shown in Figure 3, activity is highest at 3pm, which is the same as during the week. The sharpest increase in activity occurs between the hours of 8am and 9am, which is the same as during the work week as well. Female work activity on weekends is much flatter across the day, in contrast to the pattern during work days. As shown in Table 2, between 9am and 11pm, male work increases by 45% compared to the night hours. Female work increases by 19% less than this, or alternatively, 26%.

One explanation for this flat pattern for women on the weekends is that women are working less than men on these days. In Figure 4, I plot the coefficients of the regression of work activity by day of week. Activity by men and women decreases on Saturday and Sunday, but decreases more for users identified as female. As reported in Table 3, female work activity decreases by an additional 27% and 26% on Saturday and Sunday over the male decrease in activity. This is roughly equivalent to the magnitude of the difference between men and women at night.

3.3 Conclusion on Work Patterns

Worker surveys show that women, and especially women with children, work fewer hours per week than non-mothers even when fully employed. Using observational data, the above sections show when in the work week men are working when women are not. Women work a more traditional Monday through Friday, 9am to 5pm week than men in this data.

4 Gendered Reactions to Childcare Shocks

There a variety of reasons why women may choose to schedule their work hours differently than men. Women may have a greater taste for leisure time on nights and weekends than men do. Women could also be engaged in informal labor that is time-sensitive. In this section, I consider childcare availability as a key explanation for this variation. Formal childcare is less available on nights and weekends, and so the comparative drop in activity on nights and weekends may be because women are taking care of children more than men are. In order to separate childcare from other mechanisms, I evaluate unexpected interruptions in the school calendar.

Identification Strategy. I utilize unexpected breaks in the school calendar in order to identify how men and women respond to childcare availability. I specifically look at school

closures due to snow and other winter weather conditions. Snow days are helpful in two primary ways. First, snow closures are unexpected and so individuals are not able to plan for alternative childcare options in advance. Second, severe weather conditions leading to school closures typically cause a cancellation in daycare and can impact the ability of at-home childcare workers (nannies, babysitters, etc.) to commute to the individual’s home or vice versa. This means that most childcare options will be unavailable to parents,

For my analysis, I use the following general specification, which compares the activity on a winter weather school closure day to activity on a normal day for both men and women.

$$\text{Log}(\text{Activity}_{tgd}) = \beta_0 + \theta \text{Female}_g + \sigma \text{Snowday}_{td} + \alpha \text{Snowday}_{td} \times \text{Female}_g + \sum_{d=1}^7 \delta_d \text{Day}_t + \epsilon_{tgd}$$

Activity is a general term that refers to GitHub push and commit events as described in Section 1. This is the total amount of activity on date t for users of gender g on day of the week d . I include a fixed effect for users identified as female, *Female*, in order to control for baseline differences in activity levels. *Snowday* is a indicator variable for winter weather school closures. The attached coefficient σ is the impact of one of these school closures on the activity of the baseline group. The baseline group in the specification is the set of users identified as male. The coefficient α captures the differential effect of a school closure on female activity. I include fixed effects for the day of the week, δ_d .

Both α and σ are parameters of interest, as they capture the impact of a snowday on men and women. This paper is specifically interested in the differential impact of a childcare shock on women, α . I allow for errors with heteroskedasticity and auto-correlation by using Newey-West estimation and allowing for a two week lag.⁶

Results. Table 4 summarizes the results of this analysis. My preferred specification is shown in Column (4). The estimated impact of a school closure on female-identified users is consistent across the four columns, although statistical significance does vary. In all four

⁶Estimations with robust standard errors produce similar results.

specifications, the impact of a school closure on male users is not statistically significant. In my preferred specification, a snow day decreases the amount of activity by female-identified users by 34%. In my preferred specification, I include day of week, month of year, female interacted with day of week, and day of week interacted with month of year fixed effects.

4.1 Threats to Identification

In the following section, I more carefully explore two aspects of the childcare causality story I am considering.

It is possible that I am identifying a change in github activity, but that this activity is hobby- and not work-related. In order to evaluate this, I conduct my analysis again, but limit the analysis to the subset of activity that is explicitly associated with top tech companies.⁷

In the second part of this section, I proxy for parental likelihood by controlling for the estimated age of the users. In order to more carefully identify female users who are likely to have parental obligations, I consider the effect on women between the ages of thirty and forty. Women in this age group are more likely to have young children in their household (Pew Research Center, 2015). As described in Section 1, I connect LinkedIn and GitHub user information to construct a variable with the approximate age of these users.

Company-Owned Projects. In GitHub, I am able to identify if user activity is associated with a project that is owned by a major tech company. I conduct my main analysis again using the subsample of activity that is associated with these companies. This analysis is reported in Table 5. When we examine activity that is explicitly related to a top tech company, the magnitude of the decrease in activity by women is larger. In this subsample, the decrease in activity is 66%, whereas the larger sample shows a 34% decrease.

⁷These companies are as follows: Amazon, Comcast, Facebook, Google, Intel, IBM, Microsoft, Red Hat, and SAP.

Estimated Age. In Table 6, I report the estimated effect of a school closure on female activity when considering women between the ages of 30 and 40. When an indicator variable for women between these ages is included, the overall effect for women loses statistical significance. Instead, the estimated effect on women between 30 and 40 is statistically significant and much larger at -80%. This suggests that the decrease in work activity for women overall is being driven caused by an increase in childcare responsibilities, as women in this age group are more likely to have young children.

5 Robustness

I consider a couple robustness checks of my primary specification. I construct a sample of placebo female users and create placebo school closure days to examine if I find statistical significance in either situation. In both situations, there is no statistical significance. Additionally, I consider an alternate specification style where the unit of observation is an individual user.

Random Female Assignment to Male Coders. I remove all users identified as female from the sample. In this new sample, I randomly assign 11 % of users to a “female” group and recreate my primary school closure specification. In the true sample, 11% of users are identified in female. In this analysis, as reported in Table 8, I find no statistical significance.

Placebo Snowdays. I construct a set of placebo school closures by moving the true school closures back one year in time. I then conduct my main analysis on this new set of closures to see if I find statistical significance. In Table 9, I report the estimates. There is no statistical significance for any of the coefficients.

Individual Analysis. In Table 7, I consider an alternate specification style. In previous analyses, the unit of observation is the aggregate activity for a given gender. In this table,

the unit of observation is the individual user. By doing this, I am able to include individual fixed effects that account for differences in GitHub interaction style. By doing so, I can control for compositional changes in the users who are active on any given day.

In Column (1), I repeat the preferred column from Table 4, the main specification of this paper. In Column (2), I consider the model with individual observations and fixed effects. In this column, all users are included. The decrease for women is statistically significant, but much smaller in magnitude. In this data, there are many users who interact with the platform very infrequently which causes the log specification to be less than ideal. In Column (3), I repeat this analysis, but only keep users who are present in the data at least 25% of work days. In Column (4), this is increased to 50% of work days. As I increase the frequency that users must be present in the data, we approach the magnitude of the decrease found in the aggregate analysis.

6 Discussion

In order to properly consider the welfare and policy implications of these findings, we need to formalize our goal and consider the different possible solutions in turn. Additionally, it is possible that if the difference in earnings is not coming from gender-based discrimination, society may choose to accept the resulting wage difference as is. If policy does seek to change the status quo, the first key consideration is whether policy should seek to equalize the earnings of men and women in the workforce or to equalize the return to hours of men and women in the work force. In this section, I consider the ways in which earnings and hourly earnings could be equalized.

Equalizing Overall Earnings. If we would like men and women to earn the same amount of money on average, interventions aimed at equalizing male and female work hours should

be part of policy interventions.⁸ This could be achieved by increasing female work hours or decreasing male work hours. In this data, and as reported in worker surveys (Cha and Weeden, 2014; Noonan et al., 2005), women report working at least 40 hours a week on average. In order to equalize hours, women could work more on nights and weekends, on average, or men could work less during these times.

As identified in the second part of the paper, female work activity is more responsive to childcare considerations. In order to increase female work hours, we could increase male childcare responsibilities within households with male and female co-parents. An extensive literature examines intra-household bargaining. For a thorough literature review, see Bobonis, 2009. These decisions are heavily impacted by culture and have been shown to be resistant to interventions. Although, recent work suggests that paternal leave policies may have long-lasting impacts on the participation of men in child-rearing (Buenning, 2015).

Alternatively, policy interventions could be designed which would decrease childcare responsibilities for the household overall and therefore free up hours for women to devote to formal labor. Given that the literature, this work included, suggests that women are working at least forty hours a week, decreasing within-household childcare responsibilities would mean increasing outside-household childcare to greater than forty hours a week. While this is possible, it is not obvious that this would be welfare-improving for parents or children.

In order to decrease male work hours, another possible solution is implementing a maximum weekly work hour. Given that the difference is primarily driven by work on nights and weekends, removing the possibility for night and weekend work could significantly reduce to difference in work activity between men and women. Although theoretically sound, many professions do not formally track employee work hours or have systems that would allow for employees to underreport their hours. Individual workers would have strong incentive to misreport their works hours in order to receive higher compensation for higher output.

⁸It could be possible for women to make the same amount of money while still working fewer hours, but this would require women to make more money per hour on average. Given that women currently make less per hour (Blau and Katz, 2017), this goal would be a stronger version of seeking to *equalize* earnings per hour.

Given that the gender wage gap is largest in the top of the distribution, which is also where this underreporting is most possible, an intervention of this nature would most likely be ineffectual.

Equalizing Returns to Work Hours. Instead of aiming to equalize overall earnings between men and women, the focus could be on equalizing the earnings per hour. A recent theory has emerged that a key component of the wage gap between men and women is the exponential returns to long work hours. Professions with the highest gender pay parity are those with more linear relationships between hours input and wage (Goldin and Katz, 2016; Claire Cain Miller, 2019). In these professions, women are able to scale down their work hours without outsized decreases in their salary.

Interventions aimed at making the relationship between hours input and wage more linear could be undertaken by individual firms and industries without the need for government intervention. These movements could give them a competitive advantage at recruiting top female talent. In pharmacy, technological changes allowed individual pharmacists to substitute work between each other with little cost to productivity (Goldin and Katz, 2016). These advances allow individual workers to scale their work hours up or down without causing inefficiencies to the firm. As such, women are able to engage in fewer work hours without losing out on outsized salary returns.

Bunching At 40 Hours. One plausible explanation for the 9am to 5pm work day of female workers is that women are working as much as possible subject to the constraint of childcare availability. Another plausible explanation is that women are working as few hours as possible subject to the constraint of maintaining employment. In this light, the exogenous shock of a public school winter weather school closure provides a need for childcare within the household, but also a reasonable explanation within the workforce for why the female employee is not working that day.

In the earlier policy discussions, I consider how to increase female work hours. Another

important question is whether women could be better served by increasing opportunities for part-time work. As discussed above, creating a more linear relationship between wages and hours input would decrease the disincentive to scale hours down. Additionally, if labor was more substitutable between workers, the firm might be more amenable to these choices since they would not face inefficiencies from having multiple employees as opposed to one.

7 Conclusion

For women, there is an increase in happiness associated with having a family, and an increase in happiness associated with having a career (Bertrand, 2013). Unfortunately, these two increases do not add together to lead to a happiness premium for women with families *and* careers (Bertrand, 2013). In response to the demanding nature of raising children while working, we see women exit the labor force in their 30s and 40s, re-enter as children age, and postpone retirement (Goldin & Mitchell, 2017). For women who continue in the labor market, there is a motherhood penalty that has remained near constant in magnitude since 1986 (Jee, Misra, & Murray-Close, 2018).

In this work, I first document that the female work week follows a more traditional pattern than the male work week. For women as compared to men, we see that work is concentrated between 9am to 5am and that work decreases more on the weekends. Survey work identified that women report working fewer hours than men, but to my knowledge this is the first work corroborating this information. Additionally, I identify *where* in the course of the week these hours are coming from. I use a natural experiment of unplanned public school closures due to winter weather to demonstrate that changes in childcare availability impact female work activity, but not male. This second analysis is consistent with the story that appears in the first half of the paper — women are not working when children are not in school, but men still are.

In this work, I consider the impact of school closures on female work activity. The

implications are potentially much broader. There are many situations in which households may be faced with an increase in childcare responsibilities. These circumstances may be idiosyncratic to the individual household, as in the case of a medical issue for example, and they may also appear during larger shocks. The results in this paper suggest that school closures during COVID19 are likely to impact female work activity much more than male activity. Statistics coming out of the United States Bureau of Labor and Statistics confirm this. In September 2020, approximately 78,000 men exited the labor force and 617,000 women (U.S. Bureau of Labor Statistics, 2020). With a team of co-authors, I am using a similar estimation strategy as in this paper to estimate the precise gender differences in work response to this pandemic.

Broadly, life and parenthood are rife with shocks that demand increased household labor. The results from this paper suggest that these shocks will impact female workforce activity much more than male. Over a career, these differences may lead to broad differences in career trajectories and compensation.

References

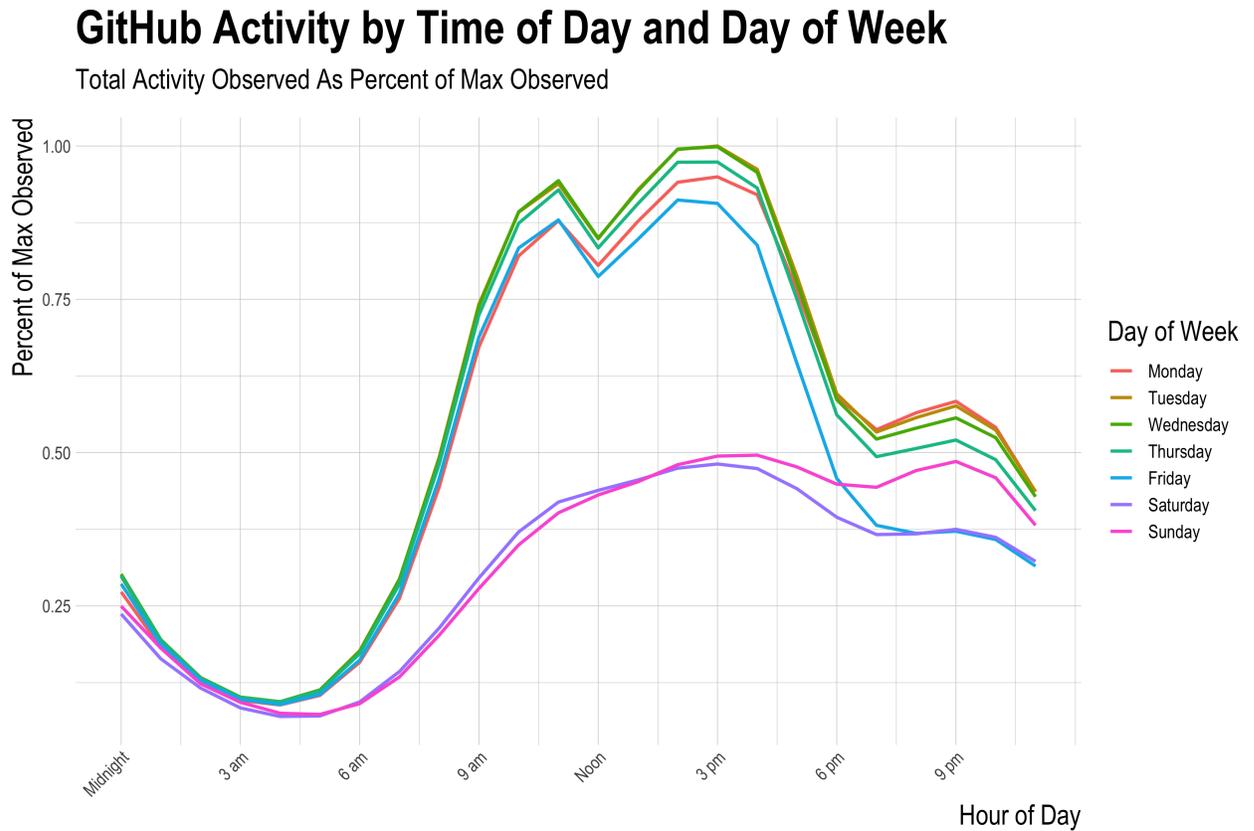
- American Association of University Women. (2019). Women’s Student Debt Crisis in the United States. *American Association of University Women*.
- Bertrand, M. (2013). Career, Family, and the Well-Being of College-Educated Women. *American Economic Review: Papers & Proceedings*, 103(3), 244–250.
- Bertrand, M., Goldin, C., & Katz, L. F. (2010). Dynamic of the Gender Gap for Young Professionals in the Financial and Corporate Sectors. *American Economic Journal: Applied Economics*, 2(3), 228–255.
- Blau, F. D. & Kahn, L. M. (2013). Female Labor Supply: Why Is the United States Falling Behind. *American Economic Review: Papers & Proceedings*, 103(3), 251–256.
- Blau, F. D. & Katz, L. F. (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature*, 55(3), 789–865.
- Blevins, C. & Mullen, L. (2015). Jane, John ... Leslie? A Historical Method of Algorithmic Gender Prediction. *Digital Humanities Quarterly*, 9(3).
- Bobonis, G. J. (2009). Is the Allocation of Resources within the Household Efficient? New Evidence from a Randomized Experiment. *Journal of Political Economy*.
- Buenning, M. (2015). What Happens after the ‘Daddy Months’? Fathers’ Involvement in Paid Work, Childcare, and Housework after Taking Parental Leave in Germany. *European Sociological Review*.
- Cha, Y. & Weeden, K. A. (2014). Overwork and the Slow Convergence in the Gender Gap in Wages. *American Sociological Review*, 79(3), 457–484.
- Claire Cain Miller. (2019, August 21). How Medicine Became the Stealth Family-Friendly Profession. *New York Times*. Retrieved from <https://www.nytimes.com/2019/08/21/upshot/medicine-family-friendly-profession-women.html>
- GitHub. (2019). The State of the Octoverse. Retrieved from <https://octoverse.github.com>
- Goldin, C. & Katz, L. F. (2016). A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation. *Journal of Labor Economics*, 34(5), 705–746.

- Goldin, C., Kerr, S. P., Olivetti, C., & Barth, E. (2017). The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census. *American Economic Review: Papers & Proceedings*, 107(5), 110–114.
- Goldin, C. & Mitchell, J. (2017). The New Life Cycle of Women’s Employment: Disappearing Humps, Sagging Middles, Expanding Tops. *Journal of Economic Perspectives*, 31(1), 161–182.
- Jee, E., Misra, J., & Murray-Close, M. (2018). Motherhood penalties in the U.S., 2986-2014. *Working paper*. Retrieved from <https://equitablegrowth.org/working-papers/motherhood-penalties/>
- Joint Economic Committee. (2016). Gender Pay Inequality: Consequences for Women, Families and the Economy. *United States Congress*. Retrieved from https://www.jec.senate.gov/public/_cache/files/0779dc2f-4a4e-4386-b847-9ae919735acc/gender-pay-inequality----us-congress-joint-economic-committee.pdf
- Miller, C. C. (2015a, July 30). Millennial Men Aren’t the Dads They Thought They’d Be. *The New York Times*. Retrieved from <https://www.nytimes.com/2015/07/31/upshot/millennial-men-find-work-and-family-hard-to-balance.html>
- Miller, C. C. (2015b, May 28). The 24/7 Work Culture’s Toll on Families and Gender Equality. *The New York Times*. Retrieved from <https://www.nytimes.com/2015/05/31/upshot/the-24-7-work-cultures-toll-on-families-and-gender-equality.html>
- Miller, C. C. (2019, April). Women Did Everything Right. Then Work Got ‘Greedy.’ *The New York Times*. Retrieved from <https://www.nytimes.com/2019/04/26/upshot/women-long-hours-greedy-professions.html>
- National Center for Education Statistics. (2017). Bachelor’s degrees conferred by post-secondary institutions, by race/ethnicity and sex of student: Selected years, 1976-77 through 2015-2016. *Digest of Education Statistics*. Retrieved from https://nces.ed.gov/programs/digest/d17/tables/dt17_322.20.asp

- Noonan, M. C., Corcoran, M. E., & Courant, P. N. (2005). Pay Differences Among the Highly Trained: Cohort Differences in the Sex Gap in Lawyers' Earnings. *Social Forces*, 84(2), 851–870.
- Pew Research Center. (2015). *For most highly educated women, motherhood doesn't start until the 30s*. Retrieved from <https://www.pewresearch.org/fact-tank/2015/01/15/most-highly-educated-women-motherhood-doesnt-start-until-the-30s/>
- Speer, J. D. (2019). Bye Bye Ms. American Sci: Women and the Leaky STEM Pipeline. *Working paper*. Retrieved from <https://sites.google.com/site/jaminspeer/research>
- United States Court of Appeals for the Ninth Circuit. (2019). *hiQ Labs, Inc. v. LinkedIn Corporation*. No. 17-16783. September 9, 2019.
- U.S. Bureau of Labor Statistics. (2020). *Employment Situation Summary – September 2020*.
- U.S. Census Bureau. (2017a). Occupation by sex and median earnings in the past 12 months. *American Fact Finder*. Retrieved from <https://factfinder.census.gov/>
- U.S. Census Bureau. (2017b). Poverty Status in the Past 12 Months by Sex and Age. *American Community Survey*. Retrieved from <https://data.census.gov/cedsci/>
- van den Hurk, A., Meelissen, M., & van Langen, A. (2019). Interventions in education to prevent STEM pipeline leakage. *International Journal of Science Education*, 41(2), 150–164.
- Weeden, K. A., Cha, Y., & Bucca, M. (2016). Long Work Hours, Part-Time Work, and Trends in the Gender Gap in Pay, the Motherhood Wage Penalty, and the Fatherhood Wage Premium. *The Russel Sage Foundation Journal of the Social Sciences*, 2(4), 71–102.

Figures and Tables

Figure 1: Plot of GitHub Activity by Time of Day and Day of Week

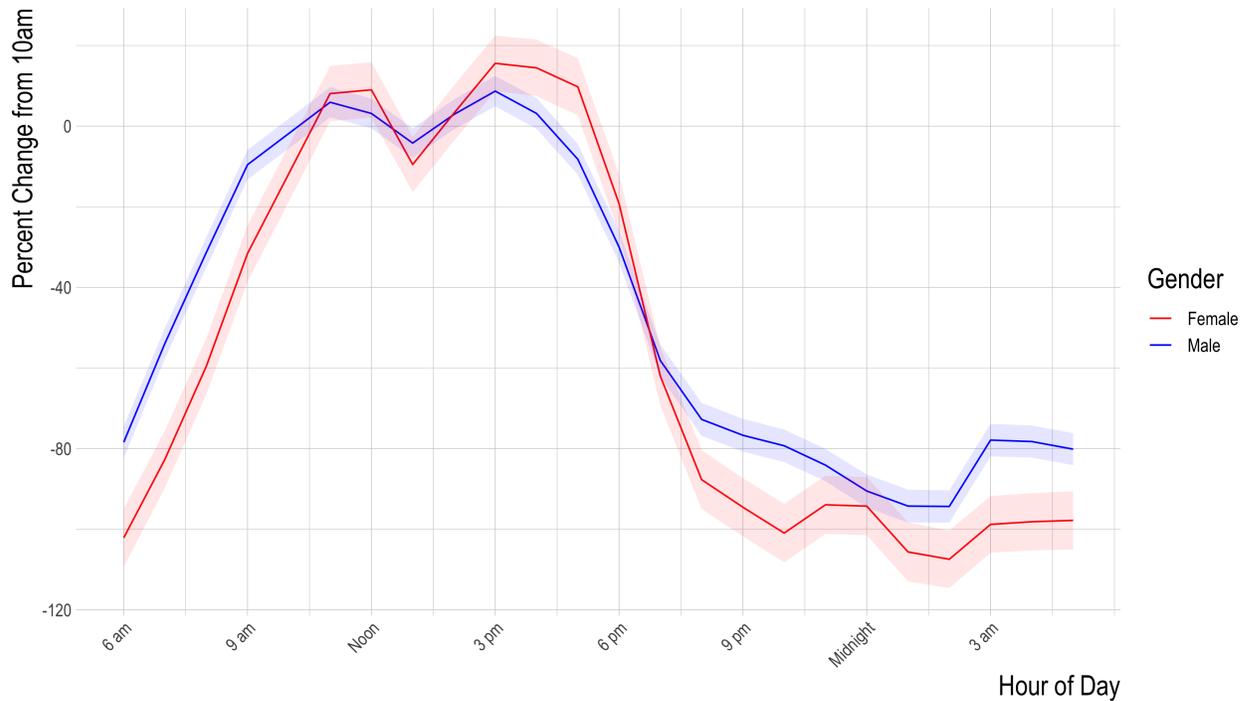


This figure plots the raw data. The unit of observation is the total amount of activity observed in a day of week, hour of day combination. These totals have been divided by the max observation. As such, y-axis reports these totals as the 'percent of maximum activity observed'.

Figure 2: Plot of Work Activity by Gender on Weekdays

Daily Work Cycle for Men and Women

Plot of Coefficients from Regression on the Log of Total Hourly Activity

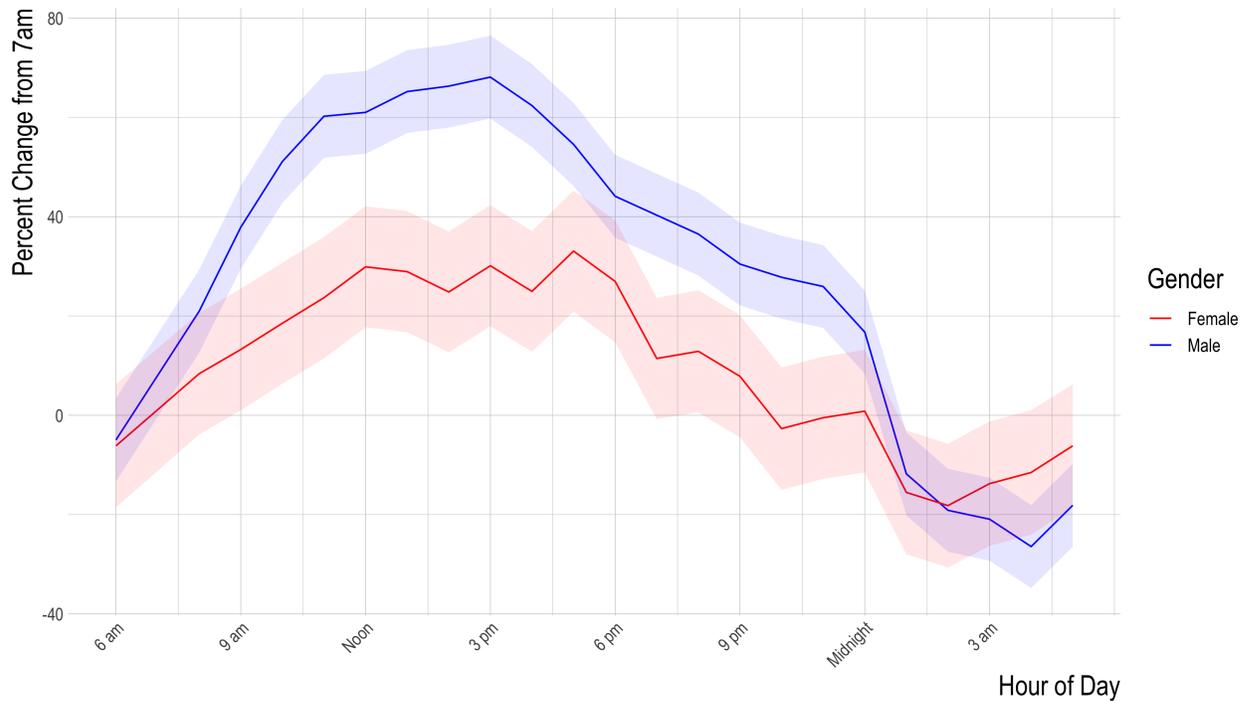


This figure plots the coefficients of a regression on the log of total activity observed in an hour. An observation is the total activity by gender for each date-hour. Fixed effects for gender, day of week, day of week interacted with gender, month of year, and month of year interacted with day of week have been removed. The 95% confidence intervals for the point estimates have been included as bands.

Figure 3: Plot of Work Activity by Gender on Weekends

Weekend Day Work Cycle for Men and Women

Plot of Coefficients from Regression on the Log of Total Hourly Activity

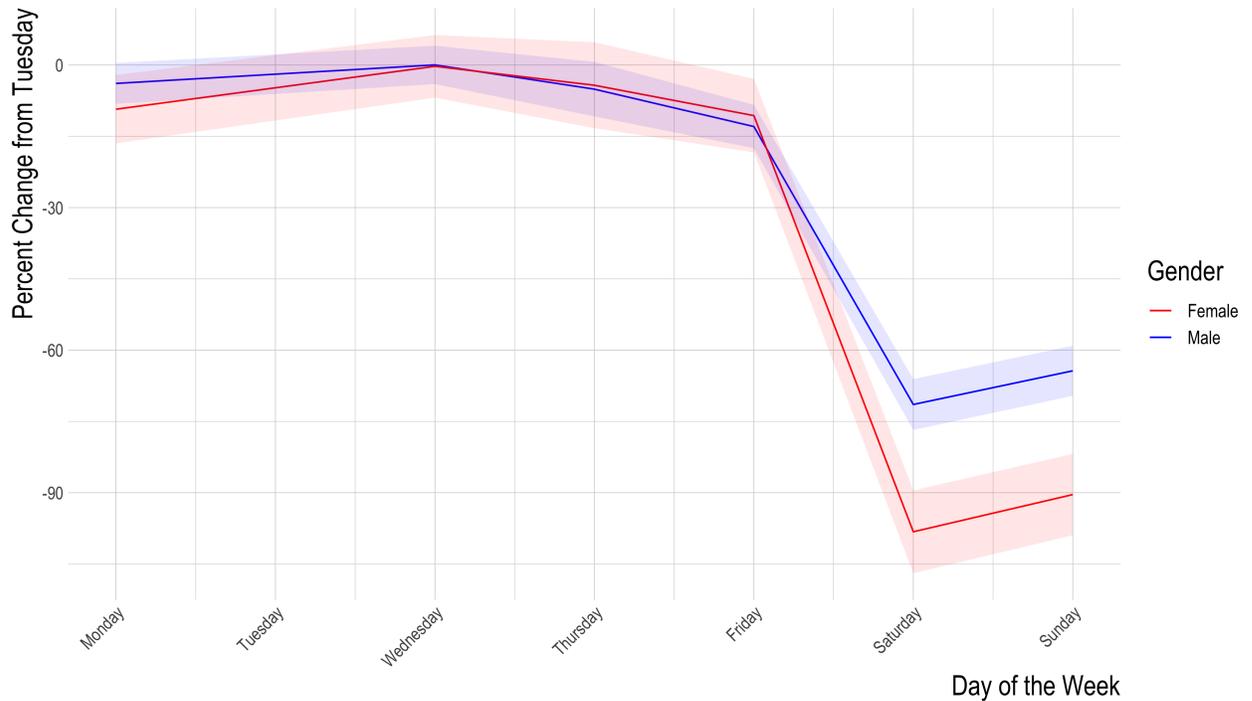


This figure plots the coefficients of a regression on the log of total activity observed in an hour. An observation is the total activity by gender for each date-hour where the date is on a Saturday or Sunday. Fixed effects for gender, day of week, day of week interacted with gender, month of year, and month of year interacted with day of week have been removed. The 95% confidence intervals for the point estimates have been included as bands.

Figure 4: Plot of Work Activity by Gender Across the Week

Weekly Work Cycle for Men and Women

Plot of Coefficients from Regression on the Log of Total Daily Activity



This figure plots the coefficients of a regression on the log of total activity observed in a day. An observation is the total activity by gender for each date. Fixed effects for gender, month of year, and month of year interacted with gender have been removed. The 95% confidence intervals for the point estimates have been included as bands.

Table 1: **Difference in Workday Between Men and Women**

VARIABLES	(1)	(2)	(3)	(4)
Female × 2am to 8am	-0.272*** (0.024)	-0.272*** (0.024)	-0.272*** (0.024)	-0.272*** (0.024)
9am to 3pm	0.019 (0.021)	0.019 (0.021)	0.019 (0.021)	0.019 (0.021)
4pm to 5pm	0.109*** (0.027)	0.109*** (0.027)	0.109*** (0.027)	0.109*** (0.027)
6pm to 7pm	-0.283*** (0.032)	-0.283*** (0.032)	-0.283*** (0.032)	-0.283*** (0.032)
8pm to 10pm	-0.193*** (0.033)	-0.193*** (0.033)	-0.193*** (0.033)	-0.193*** (0.033)
11pm to 2am	-0.226*** (0.031)	-0.226*** (0.031)	-0.226*** (0.031)	-0.226*** (0.031)
Observations	50020	50020	50020	50020
R-squared	0.743	0.743	0.753	0.754
Day of Week	Yes	Yes	Yes	Yes
Female × Day of Week	No	Yes	Yes	Yes
Month	No	No	Yes	Yes
Day of Week × Month	No	No	No	Yes

*** p<0.01, ** p<0.05, * p<0.1

This table reports the coefficients from a regression of the log of activity on the hour of the day. An observation is the total activity for a date-hour combination for each gender. This table includes observations from weekdays only. The coefficients are interpreted as percent changes from the baseline hour of 10am. In all specifications, gender is controlled for. The reported coefficients above are the differential change for women for the specified time period. E.g. Female activity decreases 19.3% more than male work activity during 6pm to 7pm as compared to 10am work activity. The reported errors are Newey-West with a two week lag.

Table 2: **Difference in Weekend Day Between Men and Women**

VARIABLES	(1)	(2)	(3)	(4)
Female × 9am to 11pm	-0.185*** (0.046)	-0.185*** (0.046)	-0.185*** (0.046)	-0.185*** (0.046)
9am to 11pm	0.457*** (0.021)	0.457*** (0.021)	0.457*** (0.021)	0.457*** (0.021)
Observations	20056	20056	20056	20056
R-squared	0.759	0.759	0.768	0.769
Day of Week	Yes	Yes	Yes	Yes
Female × Day of Week	No	Yes	Yes	Yes
Month	No	No	Yes	Yes
Day of Week × Month	No	No	No	Yes

*** p<0.01, ** p<0.05, * p<0.1

This table reports the coefficients from a regression of the log of activity for each hour of the day. An observation is the total activity for a date-hour combination for each gender. This table includes observations from Saturday and Sunday only. The coefficients are interpreted as percent changes from the baseline hour of 7am. In all specifications, gender is controlled for. The reported coefficients above are the differential change for woman for the specified time period. E.g. Female activity decreases 18.5% more than male work activity from 9am to 11pm. The reported errors are Newey-West with a two week lag.

Table 3: **Difference in Week Between Men and Women**

VARIABLES	(1)	(2)	(3)
Female × Monday	-0.054 (0.078)	-0.054 (0.074)	-0.055 (0.074)
Wednesday	-0.003 (0.074)	-0.003 (0.071)	-0.003 (0.071)
Thursday	0.008 (0.079)	0.008 (0.076)	0.007 (0.076)
Friday	0.023 (0.075)	0.023 (0.072)	0.022 (0.072)
Saturday	-0.268*** (0.072)	-0.268*** (0.069)	-0.270*** (0.069)
Sunday	-0.260*** (0.071)	-0.260*** (0.068)	-0.262*** (0.068)
Observations	2922	2922	2922
R-squared	0.841	0.853	0.853
Month	No	Yes	Yes
Female × Month	No	No	Yes

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

This table reports the coefficients from a regression of the log of activity. An observation is the total activity for a date for each gender. The coefficients are interpreted as percent changes in the amount of work from the baseline day of Tuesday. In all specifications, gender is controlled for. The reported coefficients above are the differential change for woman for the specified time period. E.g. Female activity decreases 27% more than male work activity on Saturday as compared to Tuesday. The reported errors are Newey-West with a two week lag.

Table 4: **Difference Between Men and Women in Public School Snowday Response**

VARIABLES	(1)	(2)	(3)	(4)
Female \times Snowday	-0.310 (0.192)	-0.343* (0.192)	-0.343*** (0.115)	-0.343*** (0.112)
Snowday	0.115 (0.136)	0.132 (0.136)	0.008 (0.084)	0.010 (0.084)
Observations	1032	1032	1032	1032
R-squared	0.917	0.917	0.971	0.974
Day of Week	Yes	Yes	Yes	Yes
Female \times Day of Week	No	Yes	Yes	Yes
Month	No	No	Yes	Yes
Day of Week \times Month	No	No	No	Yes

*** p<0.01, ** p<0.05, * p<0.1

This table reports the coefficients from a regression of the log of activity. An observation is the total activity for a date for each gender. The coefficients are interpreted as percent changes from a non-snowday. In all specifications, gender is controlled for. The reported coefficients above are the differential change for woman on a snowday and the change overall. The reported errors are Newey-West with a two week lag.

Table 5: **Snowday Response, Explicitly Company-Related Activity**

VARIABLES	(1)	(2)	(3)	(4)
Female \times Snowday	-0.541* (0.306)	-0.666** (0.304)	-0.666** (0.299)	-0.666** (0.293)
Snowday	0.105 (0.217)	0.167 (0.215)	0.056 (0.219)	0.154 (0.219)
Observations	914	914	914	914
R-squared	0.865	0.868	0.874	0.888
Day of Week	Yes	Yes	Yes	Yes
Female \times Day of Week	No	Yes	Yes	Yes
Month	No	No	Yes	Yes
Day of Week \times Month	No	No	No	Yes

*** p<0.01, ** p<0.05, * p<0.1

This table reports the coefficients from a regression on the log of activity. An observation is the total activity for a date for each gender. This analysis only includes activity that is explicitly associated with a major Seattle technology company. The coefficients are interpreted as percent changes from a non-snowday. In all specifications, gender is controlled for. The reported coefficients above are the differential change for woman on a snowday and the change overall. The reported errors are Newey-West with a two week lag.

Table 6: Snowday Response, Estimated Age Included

VARIABLES	(1)	(2)	(3)	(4)
Female × Snowday	-0.377*** (0.146)	-0.126 (0.145)	-0.125 (0.145)	-0.125 (0.140)
Female × Thirty to Forty × Snowday		-0.794*** (0.226)	-0.795*** (0.226)	-0.794*** (0.213)
Observations	972	2834	2834	2834
R-squared	0.904	0.876	0.881	0.885
Day of Week	Yes	Yes	Yes	Yes
Female × Day of Week	Yes	Yes	Yes	Yes
Month	Yes	No	Yes	Yes
Day of Week × Month	No	No	No	Yes

*** p<0.01, ** p<0.05, * p<0.1

This table reports the coefficients from a regression on the log of activity. An observation is the total activity for a date for each gender-age group. This analysis includes two age groups, individuals who are estimated to be between thirty and forty in one group and all others in the other. This age group is used to identify individuals who are likely to be parents of younger children. The coefficients are interpreted as percent changes from a non-snowday. In all specifications, gender is controlled for. The reported errors are Newey-West with a two week lag.

Table 7: Snowday Response, Individual Analysis

VARIABLES	(1)	(2)	(3)	(4)
	<i>Aggregate</i>	<i>Individual</i>	<i>Individual 25% of Days</i>	<i>Individual 50% of Days</i>
Female × Snowday	-0.343*** (0.112)	-0.072*** (0.026)	-0.173*** (0.061)	-0.214** (0.102)
Snowday	0.010 (0.084)	-0.028*** (0.007)	-0.027 (0.017)	-0.003 (0.030)
Observations	1032	2.259e+06	913836	435504
R-squared	0.974	0.023	0.050	0.083
Day of Week	Yes	Yes	Yes	Yes
Female × Day of Week	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
Day of Week × Month	Yes	Yes	Yes	Yes
Number of Individuals		4,378	1,771	844

*** p<0.01, ** p<0.05, * p<0.1

This table reports the coefficients from a regression of the log of activity where the unit of observation is varied. In Column (1) the preferred specification in my analysis is reported where a unit of observation is the total activity on a date by gender. In Columns (2)-(4), a unit of observation is an individual's total activity on a date. In Column (2), all users in the Seattle Area with identifiable gender are included. In Column (3), I limit the analysis to users who are present in the data at least 25% of work days. In Column (4), I limit analysis to users who are present in the data at least 50% of days. The reported errors are Newey-West with a two week lag.

Table 8: Snowday Response, “Placebo” Female

VARIABLES	(1)	(2)	(3)	(4)
Placebo Female \times Snowday	0.018 (0.166)	0.014 (0.167)	0.014 (0.092)	0.014 (0.088)
Snowday	0.128 (0.118)	0.130 (0.118)	-0.029 (0.068)	-0.026 (0.066)
Observations	1032	1032	1032	1032
R-squared	0.925	0.925	0.977	0.981
Day of Week	Yes	Yes	Yes	Yes
Female \times Day of Week	No	Yes	Yes	Yes
Month	No	No	Yes	Yes
Day of Week \times Month	No	No	No	Yes

*** p<0.01, ** p<0.05, * p<0.1

I construct “placebo” female observations by limiting the sample to only male users and randomly assigning 11% of these users to a new female group. This matches the percent found in the true sample. This table reports the coefficients from the regression which is run on the log of activity. An observation is the total activity for a date for each gender. The coefficients are interpreted as percent changes from a non-snowday. In all specifications, gender is controlled for. The reported coefficients above are the differential change for woman on a snowday and the change overall. The reported errors are Newey-West with a two week lag.

Table 9: Snowday Response for “Placebo” Snowdays

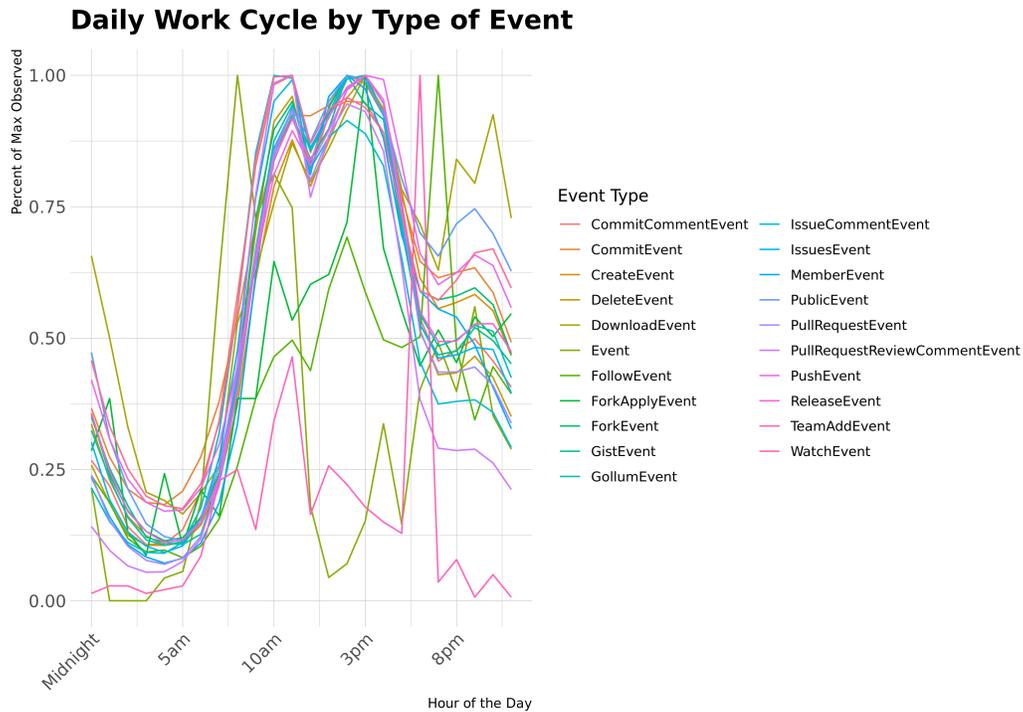
VARIABLES	(1)	(2)	(3)	(4)
Female × Placebo Snowday	0.063 (0.191)	0.066 (0.192)	0.066 (0.115)	0.066 (0.112)
Placebo Snowday	0.196 (0.135)	0.195 (0.136)	0.114 (0.084)	0.125 (0.084)
Observations	1032	1032	1032	1032
R-squared	0.917	0.917	0.971	0.974
Day of Week	Yes	Yes	Yes	Yes
Female × Day of Week	No	Yes	Yes	Yes
Month	No	No	Yes	Yes
Day of Week × Month	No	No	No	Yes

*** p<0.01, ** p<0.05, * p<0.1

I construct “placebo” snowdays by moving the snowday dates to the previous year. These are dates when Seattle Public Schools did not have public school closures. This table reports the coefficients from the regression which is run on the log of activity. An observation is the total activity for a date for each gender. The coefficients are interpreted as percent changes from a non-snowday. In all specifications, gender is controlled for. The reported coefficients above are the differential change for woman on a nowday and the change overall. The reported errors are Newey-West with a two week lag.

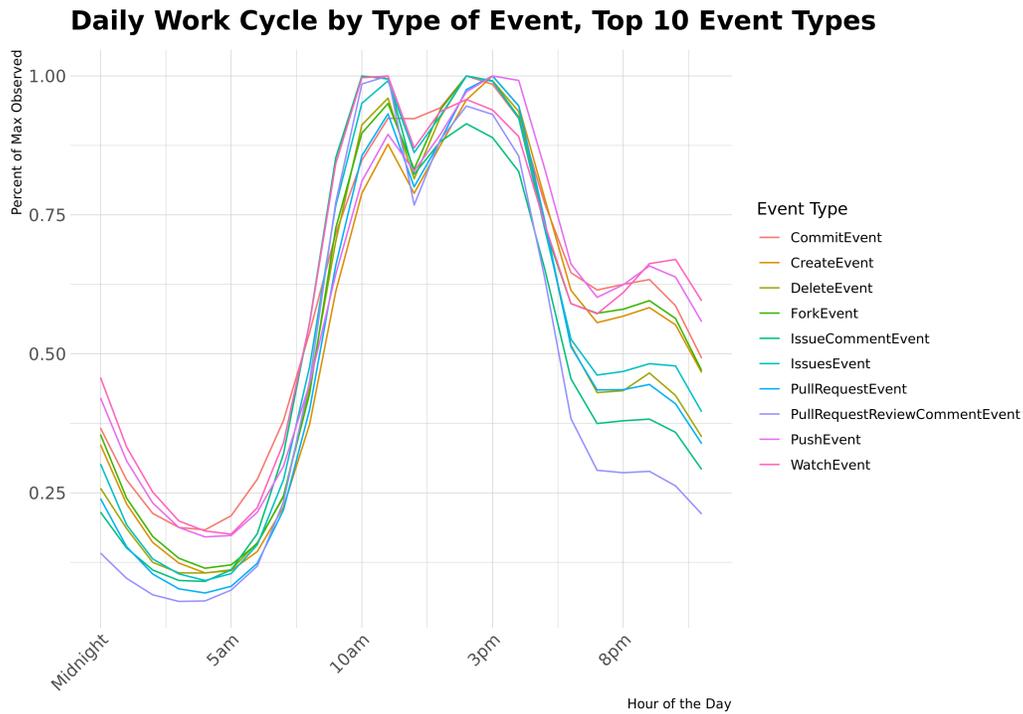
A Data Figures and Tables

Figure A1: Work Day for All Event Types



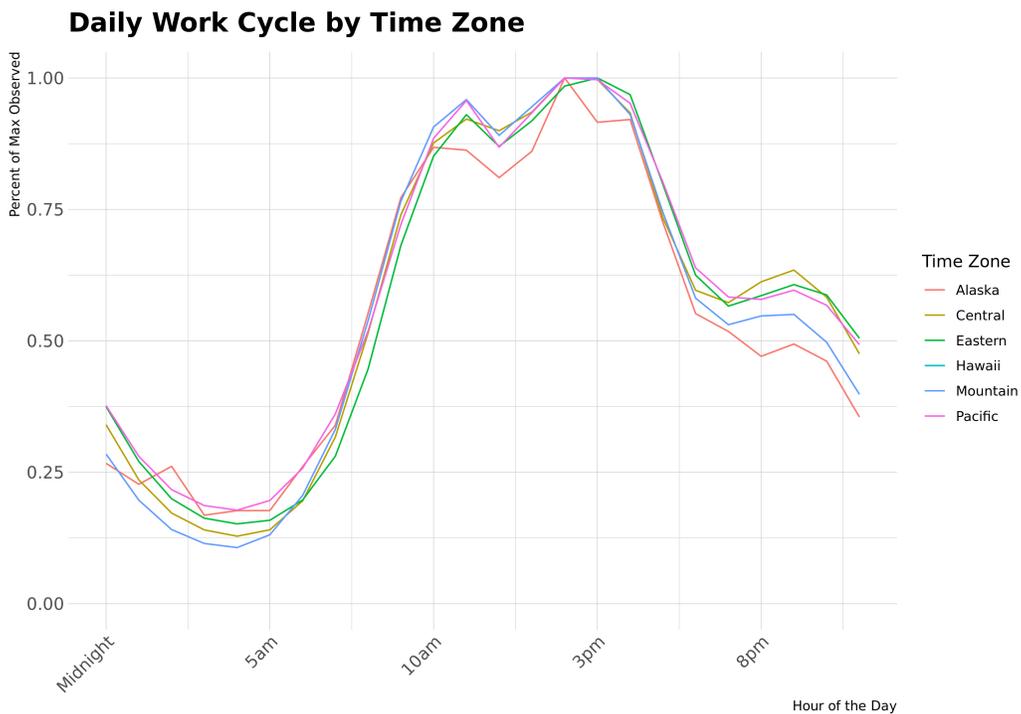
We observe all public activity in GitHub. This public activity comes in the form of different “events”. In the figure above, each line represents one type of event. The majority of event types follow the same daily cycle, but a few event types are significantly noisier and reflect a different temporal pattern. Each of these event types with a different temporal pattern represents less than 1% of the data.

Figure A2: Work Day for Common Event Types



We observe all public activity in GitHub. This public activity comes in the form of different “events”. In the figure above, each line represents one type of event. For this plot, we limit the data to types of events that represent more than 1% of observations. For these event types, the daily pattern is qualitatively similar. The greatest heterogeneity occurs in the evening. Events that involve interactions between users are less common in the evening. For the majority of my analysis, I use “commit” and “push” events. These event types follow each other closely.

Figure A3: Work Day Across United States' Time Zones



All data in GitHub Archive and GitHub Torrent is recorded in Coordinated Universal Time. I restore local time using the geographic state of the user and the time zone associated with that state. For states with multiple time zones, I use the time zone that covers the majority of the state. In the figure above, the data from each time zone is shown as an individual line. These lines reflect that data after the event times have been converted from Coordinated Universal Time to local time. The daily work pattern is similar across time zones.

Table A1: Observations by Event Type

Event Type	Percent of Observations
Commit	39.43
Push	30.49
Issue Comment	6.71
Create	6.26
Pull Request	4.17
Watch	3.97
Issues	3.14
Pull Request Review Comment	1.73
Delete	1.54
Fork	1.12
Gollum	.47
Commit Comment	.26
Release	.22
Member	.20
Follow	.12
Gist	.08
Public	.06
Event	.02
Download	.01
Fork Apply	.00
Team Add	.00

I observe public activity by all users in GitHub. This activity is described as “events”, and there are many types of events. This table documents the relative occurrence of each type of event in the data. Commit Events and Push Events are by far the most common events. Together, these events are 69.9% of the data. A commit event changes the local copy of a file. A push event changes the copy on the remote server. The remote server is shared across users who have access to a project. All event types listed above are described in detail in the Data Appendix, see Table A2.

Table A2: **Description of Event Types**

Event Type	Description
Commit	Save changes to the local copy of a file.
Push	Save changes to the remote copy of a file.
Issue Comment	An issue comment has either been created, edited, or deleted. This comment is attached to an issue that has been filed with existing code.
Create	A new branch or tag has been created. A branch is a copy of the main work that can be edited without impacting the main work. A user may choose to make the branch the main work at a later point.
Pull Request	When a user would like to contribute code to a project, this new code is issued as a pull request.
Watch	A watch event occurs when someone stars a repository. When a user stars a repository, they are choosing to follow this project.
Issues	The user identifies an issue with existing code.
Pull Request Review Comment	A pull request is being reviewed by another member of the project.
Delete	A branch or tag is deleted.
Fork	A user copies an existing project. This copy is not linked to the original project.
Gollum	Create a Wiki page.
Commit Comment	Comment on a commit that has already occurred.
Release	Release a new version of a software package.
Member	Add or remove a member of a project.
Follow	Follow the activity of another user.

Gist	Create or update a gist. A gist is a snippet of code.
Public	Switch a repository from public to private.
Download	Download a package. This event is no longer supported.
Fork Apply	Apply a patch to a fork. A patch covers the parts of someone's fork that you would like to apply to your code. This event is no longer supported.
Team Add	Add a repository to a team. A team is a group of members. This team is a subgroup of an organization.

We observe all public activity in the GitHub user platform. This activity can be one of the types listed above. All activity is described as an “event”. The left column is the event and the right column describes the event. Files are edited locally. These changes are first saved locally. These changes can then be saved to the remote server. The events listed above include actions that a user takes on the local file, actions on the remote version of the file, and actions on other users' files.